

REMARKS

A. Status of the Claims

Claim 1-22 are pending in the Application. None of the claims have been cancelled. New claims 23-39 have been added. Support for these new claims can be found throughout the specification. Therefore, claims 1-39 are at issue.

B. Rejection of Claims under 35 USC §103

Claims 1-14 and 16-22 stand rejected under 35 U.S.C. §103 as being unpatentable over U.S. Patent No. 7,031,956 to Lee in view of U.S. Patent No. 6,470,343 to O'Brien et al. Applicant respectfully submits that its invention is patentable over these references, alone and in combination.

The claimed invention relates to a database management system that includes a database schema. Claim 1 has been amended to clearly recite that Applicant's invention is directed towards a database schema. Applicant's invention is directed to a database schema. In contrast, in Lee, the metadata tables are used as a basis to allow the relational schema **22** to be *generated*, and for this reason alone a skilled person would not look to the metadata tables of Lee for the database schema management of the present invention, because the metadata tables of Lee are not a database schema.

Lee describes a mechanism for updating a relational database with supplemental data (see Abstract), which, in one example, involves adding data from an XML file into an existing database using a document type definition (DTD) (column 16, lines 17 to 20). To achieve this, the DTD **18** is loaded by the system **10** and used in metadata format to generate not manage a relational schema **22** (column 15, lines 49 to 52). Thus, data representative of the DTD is stored in metadata tables **34**, before being used by the generator **28** to generate not manage a relational schema **22** (column 15, lines 62 to 67).

Lee describes in column 18, lines 1 to 3 that the relational schema is achieved by mapping the metadata tables. This mapping requirement means that the schema would not have a similar structure to the metadata tables, otherwise such mapping would not be required.

The mapping process is actually described in some detail in Lee in column 7, lines 17 to 35. This highlights that the process includes requirements for:

(1) creating a table in the schema of the relational database corresponding to each row of the metadata item table;

(2) altering the schema of the relational database to add a column to each table in the schema corresponding to each row of the metadata attribute table related to the particular metadata item table row;

Requirement (1) therefore requires that a respective table is created for each row of the item table, which will therefore correspond to a respective table being created for each entity within the data. Similarly, requirement (2) means that there will be a column within each entity table that includes attributes (or fields) for the respective entity.

Thus, the schema of Lee uses a respective table to define each entity, with data (field values) relating to the entity being stored in the respective entity table. It is apparent from this that the methodology of Lee results in a database schema including a number of tables, each of which includes details of a single entity and corresponding fields for the entity. As will be appreciated by persons skilled in the art, this is a standard form of database schema (commonly referred to as a “table-per-object” type schema), albeit with custom tables as required by the particular data.

In contrast, the claim 1 requires:

- a) a first table to store the names of various entity types;
- b) a second table related to the first table to store the names of entities of the various entity types;
- c) a third table related to the first table to store the names of fields in respect of the various entity types;
- d) one or more value storage tables related to the second and third tables to associate stored field values with entities; and,
- e) identifiers to indicate the nature of the data to be stored in each of said tables.

The first table is used to store names of various entity types, whilst the second table stores names of entities of the specified types. Thus, in the schema of claim 1 the names of the entities are all stored in the second table, whereas in the schema of Lee a respective table is required for each entity. Consequently, the schema of Lee does not satisfy the requirement of a second table related to the first table to store the names of entities (as in Lee each table stores

only details of one entity and not plural entities), and there is also no table that stores names of entity types.

Similarly, claim 1 requires that field values are stored in the one or more value storage tables, meaning that field values associated with an entity are not stored in the same table as the entity itself, as is also the case in Lee. Consequently, the schema of Lee does not provide one or more value storage tables related to the second and third tables to associate stored field values with entities.

It is apparent that rather than use the table-per-object structure of Lee, the present claim requires an arrangement in which entities are defined by the first and second tables, with the field values being stored in separate one or more value tables. This arrangement leads to a number of advantages not taught or suggested by Lee.

For example, the schema taught by Lee is a standard table-per-object schema in which each entity is defined by a respective table within the database. Consequently, whenever it is needed to add a new entity, it is necessary to define a new table within the database. Columns must then be defined for attributes, before the table is linked or otherwise related to other database tables.

In contrast to this, by having a first table for storing entity types and a second table for storing entities, whenever a new entity is added, this can be achieved by adding one or more new entries to the existing tables, rather than be creating new tables. Similarly, field values relating to the entity can be added to existing one or more value storage tables.

Avoiding the need to add new tables significantly reduces overheads involved in updating databases to include new data. Furthermore, not only is such an arrangement not taught or suggested by Lee, but Lee in fact teaches away from this by using a schema that explicitly requires that a new table is added for each new entity. Consequently, Applicant respectfully submits the schema of Lee is relevant to the present claims.

Applicant wishes to reemphasize that in Lee the database itself is not managed in accordance with the metadata. Rather the metadata is used to define a database schema, which is in turn used to manage the database. Applicant respectfully submits that the metadata is not relevant to the present claims, which require a schema of the claimed structure.

If, however, the Examiner continues to maintain that the metadata tables of Lee could be considered as equivalent to the claimed tables, Applicant wishes to highlight that the metadata in Lee does not include one or more storage value tables to associate stored field values with entities. Instead, as stated in Lee at column 14, lines 50 to 65, the DTD is stored as metadata, whilst the data of the XML document (i.e. the content to be added to the database including the entities and the corresponding values) is extracted and stored as tables in the defined by the relational schema. Thus, the metadata stores the DTD, while the data to be stored in the database is stored in the relational schema tables, and not in the metadata.

Furthermore, the current claims require that the tables are related in a particular manner, and notably that the one or more storage value tables are related to the second and third tables. However, the metadata tables of Lee are not related to the database schema tables. Accordingly, even if the Examiner were to maintain that the claimed first, second and third tables are shown by the metadata in Lee, and that value tables are shown by the database schema tables, then the required relationship that the "*one or more value storage tables related to the second and third tables*" is not shown.

Applicant also notes that the Examiner has indicated that Lee discloses the same concerns as the current application. Lee, however, does not provide the same solution or advantages of the claimed invention. As previously stated, Lee requires a respective schema be generated for each database application. Furthermore, the schemas use a table-per-object arrangement, and consequently, when importing data schemas need to be modified by including new entity tables corresponding to new entities in the imported data.

In contrast to this, the claimed invention provides a generalised configuration that can replace existing "table-per-object" type schemas for all relational databases. This arrangement allows data, including new entities, to be added to the database without requiring the addition of new tables to an existing schema. As discussed above, this is clearly not taught or suggested by Lee.

Applicant also respectfully submits that O'Brien does not teach or suggest overcoming the deficiencies of Lee to lead to the claimed invention. O'Brien describes a schema containing tables enabling various parties to define extensions to a core data model (column 3, lines 2-5).

For example, Figure 3 of O'Brien and the accompanying descriptive text discloses a number of tables that are used for extending the data that may be stored in relation to a core table. However, it does not describe a database management system that has a schema that allows new entity types to be added to the core data model by means of generic entity type and entity tables, without requiring the addition of new tables. Applicant also notes that O'Brien also does not disclose a first table to store the names of various entity types, a second table related to the first table to store the names of instances of entities of the various entity types and further tables to store field attributes and field data, as required by claim 1. Consequently, the teaching of Lee, even when considered in the light of O'Brien, does not lead a skilled person to the teaching of claim 1, claim 18, or their respective dependent claims. Applicant respectfully submits that claims 1-14 and 16-22 are patentable over Lee in view of O'Brian.

Claim 15 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Lee and O'Brien and further in view of U.S. Patent Application Publication No. US2003/0167455 to Iborra et al. Applicant submits that claim 15 is patentable over the combination of these references.

Claim 15 depends from claim 1 and includes all of its limitation. For the reasons given above, that Lee and O'Brien do not disclose any details of the database management schema of the present invention, the addition of Iborra does not overcome the deficiencies noted in Lee and O'Brien. Therefore, claim 15 is patentable over the combination of references.

CONCLUSION

It is respectfully submitted that all of the Examiner's objections have been successfully traversed. Accordingly, it is submitted that the application is now in condition for allowance. Reconsideration and allowance of the application are courteously solicited.

Respectfully submitted,

Date: June 9, 2008

By: Monique A. Morneau
Monique A. Morneau, Reg. No. 37,893
Customer No. 1923
McDERMOTT WILL & EMERY LLP
227 West Monroe Street
Chicago, Illinois 60606-5096
(312) 372-2000

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail, postage prepaid, in an envelope addressed to: MAIL STOP RCE, Commissioner for Patents, PO Box 1450, Alexandria, VA 22313-1450, on June 9, 2008.

Sarah J. Goodnight
Sarah J. Goodnight
CHI99 4989663-1.052003.0014